MFS transcript 5

Continuity: How To Grow A Human: My Frankenstein Summer with Dr. Philip Ballepisode five Artificial Intelligence.

Siri: Hello, I am Siri. There were 17 emails received today

Dr Philip Ball: in the last episode, I talked about the Turing test, the famous test devised by mathematician, Alan Turing, to try to distinguish an artificial intelligence from a human

Siri: sunset will be at 7:28 PM today.

Dr Philip Ball: I don't think Siri is going to fool anyone into believing she is sentient

Siri: I don't have an answer for that. Is there something else I can help with?

Dr Philip Ball: All the same, even this rather crude algorithm begins to give us the feeling that we're interacting with another mind. God. Singing is harder than I thought.

And one that seems in some ways to know a lot more than we do

Siri: the square root of 1,692 is approximately 41.1339.

Dr Philip Ball: I went to Boston to explore how we're making and remaking the human body and mind in my book, how to grow it. I showed how we're already starting to grow rudimentary body parts, tissues, and organs from scratch, or rather from reprogrammed cells of our own.

As I described in earlier episodes, this might work for say the pancreas or the kidney, but growing a real working brain is much harder. Might we instead create an artificial mind, like a super advanced computer that might be capable of doing things, a human. And perhaps even feeling things a human does.

I discovered that giving AI human like powers of cognition and reason is a daunting task, not least because we don't fully understand how we learn to reason and navigate the world ourselves.

To find out where we really are with AI today and where we're headed. There could be no better place for me to visit than the research lab of computer giant. IBM based in Kendall square in Cambridge, just down the road from MIT. This is where IBM is developing Watson health. The AI system used for making medical diagnosis.

Among other things. The IBM Watson system is now being used to provide reliable information about COVID-19 from advice on symptoms and testing to public service information about transportation and unemployment benefits, as well as supercomputing resources for seeking treatments and cures.

The Kendall square center is the hub of IBM's efforts in AI, for which the company has partnered with research scientists at MIT itself in a $240 million project. I went along to speak with David Cox, who studied at Harvard and MIT at the intersection of machine intelligence and neuroscience before becoming director of the MIT IBM Watson AI.

David Cox: So we're practically part of IBM research. So I've been research is it's kind of like the last standing big industrial research lab in the bell labs kind of era. Right. But a big part of the lab really is about how do we solve the fundamental limitations of AI? Like, you know, like you've kinda said it earlier, AI is kind of dumb in a lot of ways.

Everyone's been really uncomfortable with this term artificial intelligence for a long time. 2017. And before academics, we just didn't really want to use the term artificial intelligence at all. You want to call it machine learning, deep learning if that's what we were doing. And then for whatever reason, 2018 and on we've all just, everyone's given up, we're all going to call it AI now.

And I don't know why that happened, but it happened. Google does it, we do it. Academics are doing it. And I think part of the ambition of this lab is how can we make AI genuinely intelligent? There's a, there's, there's a fundamental sense in which, you know, our AI today is very, very limited. It's not hard to find the sort of gaps our AI systems today.

Like deep learning requires huge amounts of labeled data. So you want to train a system to recognize cats or dogs. You train it on thousands and thousands of millions of images of cats, thousands of images, billions of images of dogs. And it can do that. My daughter, when she learned how to see, wonder what a cat wasn't, what a dog was.

I didn't sit there with flashcards and show her cat cat, cat, dog, dog. She saw one. She mistakenly called the dog and cat. And they were like, no, no, no, that's that's dog. And then boom, just gotten two examples. You know, she over-generalize, it it's totally fine. So why was she able to learn so fast? I don't think the answer is that somehow our brains magically need less data.

It's just where much smarter. W how we're using data. We don't require so much. Supervision is really the difference. And part of the reason is, you know, today's deep learning systems basically have to learn everything about seeing in the context of whatever task you've given them. So if you're teaching assistant cats and dogs, all it ever knows about anything about the world is through the lens of this very narrow lens of what's a cat versus what's a dog in reality.

When we look through them and my daughter saw. No. She already knew things about for legs and objects and objects can pass in front of other objects and, you know, they'll include them. And if I, you know, if I put this barrier up, when I had sight of this object behind my hand came out the other side and it wasn't still holding this phone.

Um, even, even babies who can't even verbalize yet surprise, we'll look hard at that and say, you know, realize that there's something wrong about that. So there's this whole interesting world of common sense. And the, these base-level understanding. And one of the things that's really hard about teaching today's AI methods is kind of common sense is a lot of it's, uh, unspoken and unwritten.

So if you, if you train a system on all the languages of the world, uh, it's interesting how much of that common sense isn't written anywhere? Uh, you know, we, we don't write it down because we all take it for granted. Uh, so I think if we want to build systems that can be genuinely intelligent. And for many of the applications that IBM is interested in.

Um, we're often going, we're solving, you know, like we, we solve customer's problems, but we'd really like to be able to go in and have AI systems that can go in and be much more flexible and solve the problems that don't require armies of data scientists align up the problems and then carefully annotate, huge amounts of data.

So I think it's really important that we have that base level of common sense, logic and reasoning, and being able to extend, to bring outside knowledge to the table. Because one of the reasons why my daughter was able to recognize that cat or recognize that dog, because she, she knew a lot of things. She brought a lot to the table that she had gotten through unsupervised experience with the world that she had gotten through, exploring the world for being curious about the world, you know, you have a baby and you give him a new object.

The first thing they do is they pick it up. They move it around in front of them. They might put it in their mouth. They explore. So they're constantly exploring the world. So that notion of curiosity, that notion of common sense that notion of external knowledge and building up that knowledge, that's something that I think is going to be critically important for us to kind of crack the nut and close that gap.

Looking at humans, looking at nature is a reasonable place to go looking for some of those insights, because it's basically the only existence proof for truly intelligence. That we have on earth today. And in many ways there's no guarantee that we necessarily need to look at nature or to humans to solve these problems.

But, you know, I think, I think we'd be a little bit foolish not to that's what we've built. I guess if that's the only example of the thing you're trying to build it logically, it doesn't necessarily have to be that you have to emulate it. And I think there's an easy road. Go down. That's the wrong road to say, like we have to slavishly copy every detail.

So there was a project called the human brain project in the U. That was let's simulate a brain from the, from the ion channels of the neurons up. And then magically, somehow it will do something. And that that's, that was, I think, I think I say it to say that was, uh, that was the wrong approach. Um, but at the same time, At the right level of abstraction, looking at how babies develop for instance, and what do they know and when do they know it and how do they learn that?

Um, there's a whole interesting field of child psychology that looks at these developmental milestones. When does the baby start realizing that when an object passes behind include her in, that comes out the other side, it needs to come out there. This has the same object that there's interesting work on sort of how that develops.

And I think we'd be foolish not to pay attention to that and think about how.

Dr Philip Ball: In the last episode I heard from Harvard psychologist, Toma Orman, that one route to making AI smarter could be to have it learn like a child and perhaps to make it embodied in some way so that it can develop physical intuitions about its environment.

I was struck now to hear David say that this is more or less exactly what's in the works.

David Cox: Um, so there's an interesting, uh, DARPA program that just started called machine common. We're participating in it together with her IBM or MIT kind of parts. Yeah. The basic gist is we have to have embody days. Live in an environment and learn like children, you know, and that's, that's our approach to the problem we were actually in.

And there was a requirement actually that every team have child psychologists involved in, in their teams. So it's just fascinating to get, you know, AI people and cognitive scientists and developmental people all in the same room. And it's actually interesting that there's an interesting. Labor market thing that's happening, but there's an awful lot, like neuroscientists don't get training in building AI systems, but there's a surprising number of us neuro Simon neuroscientist by training, who ended up in doing AI.

And I think it's something about the discipline of trying to understand an intelligence system and ask those questions and, and design experiments and think about the right tasks for whatever reason. I think it's not the only kind of people who are good at doing AI research, but there's a surprising number of us who've crossed over.

And are now tackling these problems. And I think it is this discipline of devoting my life to understanding this intelligence system and how it works. And now we're building artificial ones that are like, well, the idea of probing it and understanding how it works and understand what the right questions to ask.

I think. Are very compelling for people who are doing that sort of crossover from one field to another. And might

Dr Philip Ball: some of these shortcomings

David Cox: that AI

Dr Philip Ball: currently has the architectural ones? Absolutely. Okay. So

David Cox: what, what other architectures

Dr Philip Ball: could you try? Which ones do you think,

David Cox: uh, uh, needed? Uh, the architecture thing is interesting and one thing.

If, you know, a big part of what we're trying to do is actually fundamentally symbolic, right? Like a big part of what we do as humans and reason. Our ability to reason comes from our ability to manipulate symbols and language is basically a serialized stream of symbolic tokens. Neural networks was back in when I was starting out was it was a very disruptive.

Thing in the, in the nineties. So in the eighties, you know, that was going to, it was going to save the world. And then if we had an AI winter, and it was nobody's serious, worked on neural networks and all of a sudden, now the data computed caught up, but now neural networks are working. Of course there was another thread of, of AI, which was symbolic AI.

And that was also going to save the world back in the eighties. It also went through winter and nobody, it was, you know, serious work, done symbolic AI. And it hasn't enjoyed yet resurgence the way the deep networks have. But I think the interesting thing is just the same way that deep neural networks were waiting.

It was the right idea, but it wasn't tight yet because the compute and the data hadn't come yet, you know? And when we got GPU's, which sort of accidentally were the right co-processor for doing neural networks. And we had image net, which is a product of the fact that we had so many digital cameras, everyone's taking images now there's data.

They started to come into their own. I believe that symbolic AI, I think there's a lot of right ideas there too. And it's also been waiting, but what has been waiting for, I think, is neural networks. So I think there's a real, very real sense in which, uh, you know, just the same way that neural networks are waiting for the right environment, the right compute, the right ingredients.

I think now that we have neural networks, combining them with symbolic systems can be very apparent.

Dr Philip Ball: Time to explain a bit of the jargon here, symbolic AI, which was the earliest form of AI to be explored, tried to develop intelligent systems that manipulate symbols, representing things in the world, according to a set of rules, but she was, we manipulate words in language using the rules of grammar.

It explicitly maps, real concepts and objects onto the system. Neuro networks take a different approach. They are crudely like the actual architecture of the brain neuron light nodes, interconnected in a network that sends signals to one another and adjust the strength of the links between the nodes. As they learn to train a neural network, to do a task, you feed it some input data and let the strengths of the links adjust until it gives the right.

Pictures of cats producing the output cat say you don't worry about what's really going on inside the network so long as it gives the right output. Once the network is trained well enough, you can give it input data. It is never seen before, and it should generate the right output, distinguishing images of cats from dogs.

For example. This approach to ARA was explored in the 1980s, but didn't really get very far until the technique called deep learning came along in the past decade, which uses more layers of nodes in the network to get much more reliable learning most uses of AI today, such as the algorithms that beat human masters at games like chess and go use deep learning, but more and more AI researchers are realizing that deep learning has some deep floor.

It's very narrow in what it can do and can be rather easily fooled if you know how it likes common sense to keep it on track. Now, David told me some researchers think that the way forward might be a combination of both approaches, deep learning and symbolic AI.

David Cox: So we're starting to build these neuro symbolic hybrid systems where these neural networks do the things that neural networks are good at.

Like, you know, Images, all that kind of messy correlations of the world, natural language, you know, it's, it's messy. And, and you know, these, these statistical learning methods like deep, deep learning can be very powerful for doing things with that. But then when you marry them with a system, you use them to extract symbols and then you can take those symbols and do actual logical reasoning operations, uh, some pretty interesting things emerge and, and, you know, some people have taken the approach that we want to also do the symbolic processing with neural networks.

And, and that's, I wouldn't bet against that approach, but at some level, you know, we have much better symbolic computer. Structures already a colleague of mine David's DiCillo at, uh, at Google said once that, um, deep learning is crappy neural networks implemented on amazing turning machines, like our computers and symbolic reasoning and humans is crappy tournaments.

Running on amazing neural networks, right? We have this amazing, you know, biological neural network hardware. And we, we kind of barely, barely eke out the ability to like do arithmetic and things like that. And meanwhile, we have these machines that are these Silicon machines that can do amazing arithmetic and logic and reasoning and take all the world's information.

And then we're simulating neural networks on them. So there's a sense in which well. You know, if we can marry neural networks with the incredible symbolic processing power that, that digital computers have, uh, you know, maybe we don't have to like figure out how to make neural networks do that. It's about reasoning.

We just do it directly and then fuse the two systems together. Right. And, and we're, we're kind of, we're kind of marching our way down that path of saying, you know, neural networks are very powerful. We love neural networks between can we marry that with other kinds of architectures and kind of take the best of all these worlds?

Dr Philip Ball: Worth trying to make artificial

David Cox: intelligence, more

Dr Philip Ball: human, like, because we want it to do human, like tasks, like medical diagnosis or image

David Cox: recognition, voice

Dr Philip Ball: recognition, or might there be, might there be tasks that we want systems like this to do that we don't do, but we need a different sort

David Cox: of intelligence for, I mean, I suppose in some sense, neural networks do that already, but we might, there might be others,

Dr Philip Ball: you know, to what extent do we want.

Get

David Cox: developed AIS that

Dr Philip Ball: have things like

David Cox: common sense and you know, our types of abilities. And to what extent might we want to broaden that? Okay. Yeah. So, uh, the answer is yes. So, so, uh, that's a, that's, that's a, that's a great question. Uh, and there's kind of a bunch of different, a couple of pieces to it, which I think are really interesting.

So some people ask them, do you want to make a human brain in a computer and like, and have it be conscious and have emotions and all that. And my answer is. I have no reason to do that. I don't have no reason to want to do that. Like our company does not try to do that where the, like, there's no reason there's no business reason.

I don't think there's a pro social, societal reason to build a conscious computer system. Um, but things like common sense and you know, things that are closer to how we think that can. Very valuable, uh, simply because we are so effective because we have those abilities. So, you know, you think about just the task of like, I have a spreadsheet, you know, like I'm one of our client customers.

We have these like piles of spreadsheets around that have all this information and you look through it and like how you can just kind of like parachute in as a human. Read the read the data, figure out what's going on. And like you figure out what you need to do. We don't have, as systems can do that today.

And you look at the column, the, you know, the header of the column, it says, you know, factory underscore inventory. And it's like, well, okay, well to actually be able to operate in those environments, what do you need to. Well, you need to know a factory is building and inventory is like stuff that's used in the factory to build other things.

And people work in factories. There's all this stuff that very quickly now people will have our people have arms and legs and you get an injury and injury. It might be in the arm or leg. And, you know, like it's just, uh, it's actually a almost frightening amount of information we constantly bring to the table.

Even when we solve a very. Narrow specific task. And I think one of the magical things about human cognition is precisely that we have all this unspoken based level of common sense knowledge. And we just have deep, deep knowledge that seemingly irrelevant knowledge that that's actually relevant, surprisingly large amounts of the time.

If you actually want to operate in the environment with the kind of fluidity that we are able to operate in the EMR, the real way we learn that is we don't get like a core dump of a bunch of lists statements. Symbolic farm. We learned it from the environment. We learned it through experience or curiosity to investigation being embodied in the environment.

Um, but you know, the idea that we can just build like a one-off deep learning. It's only job is going to be to like categorize, you know, the parts, you know, and the factory that do this or that. I think that's going to give way to an arrow where we increasingly have as systems that know a lot of things that seem irrelevant, but they actually aren't irrelevant that we need to have that breadth of knowledge that humans have to do even simple things like that could be solved in this narrow way, but, you know, it's much more broad.

So I think in that way, we're going to move closer to how, you know, to being more human-like.

Dr Philip Ball: In the light of these developments, I've wondered what the real prospects are for making AI, that exceeds human capabilities in ways that aren't just a matter of crunching more data more quickly. And if so, would there be any need to try to give it some form of sentience or could it go on working perfectly well as a sort of zombie supermind

David Cox: now I do think, you know, we can embrace.

The capabilities that computers and computers can clearly do things that the machine that brains cannot talk about. Advance, you know, math, you know, crazy arithmetic on each numbers. Like if we, as we build these intuitive things, there's no reason not to build those increased capabilities on top. In some sense, we do that now through interfacing, like I, you know, my phone has all of the world's information on it.

I can access it, but you can imagine a much more direct connection there. So there's lots of ways we can exceed. It's not hard to imagine us being able to exceed the capabilities of humans very quickly. Once we have some of those base level things built. Also, even if we don't go to the level of human reasoning and problem solving and all that, you know, something like driving a car, you know, an autonomous vehicle, you know, even if it's not as smart as a human brain, the fact that the autonomous vehicle has cameras, you know, like you can have 14 cameras pointed in different directions.

You know, humans didn't evolve to drive cars. You know, we have these two, you know, Manipulator things that we hold on to a wheel and we have two eyes and only two eyes that facing forward. They're on this weird turn thing called a neck. We want to turn around, look around, actually turn my body around. And if I'm sick, my neck stiff, I won't do it as much in old people have trouble.

So it's, um, we're not really very well adapted for driving your car. So even assistant is a lot simpler than us. Imagine like the intelligence of a fly or intelligence of a rat. But like optimized for driving a car has got cameras pointing in all directions. It's got radar, it's got sonar. I can easily imagine getting to a point where even a simpler intelligence.

Wow. Performance humans in the domain that it's adapted for, or imagine, uh, imagine, uh, artificial intelligence that grows up and has MRI scanners for eyes, right? Like we can't process volumetric data natively. We have two dimensional present. We have to go through step slice by slice, but an AI system can exceed us that way.

So I think there's lots of opportunities. So, so in some ways it's like, I don't think we need to reach for the consciousness part and all those things. That's an undesirable. To be found over there, but on the other hand, we can reach a lot farther in other domains and just simply building things to purpose, uh, is, is, is a huge thing.

There's also no constraint on the, on the size and capacity. Our brains are, are limited by our, our ability evolutionarily to get food. You know, like brains are incredibly metabolically expensive. I'm just paying, you know, maybe a third of my energy, just having this brain and writing about 20 Watts right now.

And. That constraint was a very careful balance that evolution struck on the African Savannah, you know, eons ago. But we don't, once we figure out how to do the algorithms of intelligence, we're not going to be constrained by that anymore. So, so there will be a possibility to have intelligences that are potentially bigger than us, you know?

Cause cause we can set the constraints. Right. So intelligent, but not necessarily with anything we could call consciousness. That's right. Um, how about for doing science?

Dr Philip Ball: Obviously already, AI

David Cox: is clean learning it's

Dr Philip Ball: everywhere and there is no field of science. It doesn't have it, but. Discovery real discovery science, um,

David Cox: and you know, theory development.

Um, do you see

Dr Philip Ball: that

David Cox: feasible? That's the M we're already heading in that direction. So, so, uh, actually the team that sits like immediately right behind you is, uh, working on something called causal inference. So, you know, a lot of our AI methods today are fundamentally correlational and they have all the weaknesses come from fundamentally focusing on correlation, not causation as a new field called causal inference, which is building systems that can.

In for causes from observational data, but then also schedule tests, like basically, like, what are they, what are hypotheses that I, part of the experiments I need to do to test my hypothesis most efficiently. Uh, and that's a big part of the scientific method. So like there's an emerging increasingly this field of, you know, the fields of statistics and experimental design and AI are kind of are starting to merge.

I think that's one piece of the puzzle and piece. When you talk about discovery science, another thing that's interesting. We have a bunch of work going on. In generative modeling, where, who builds systems, the air systems. They don't just take an input and make a decision, but actually can imagine candidates for something.

So you want to design a molecule to do a particular purpose. Maybe it's a drug. Maybe it's a fit for a physical application. Um, can you build AI systems that can actually suggest, you know, things, Hey, here's something maybe you haven't thought of before. Like, why don't you try this one? You know, like this, you know, we, you know, we we've taken past examples.

You know, we're going to, we're going to be creative, like have machines that can actually have a sense of sort of creativity to explore the space more fully and kind of be your design copilot to say, you know, like maybe, maybe you should try out these drugs or what do think about this molecule. And I think that's going to be an interesting emerging thing.

You know, increasingly, you know, science is going to change because we're just going to take for granted that we have these copilots that can both do some of the hypothesis testing and thinking about how would you go about making these experiments? And people are absolutely doing that today already.

Then you're also going to have, in terms of discovery, you know, new tools that let you kind of. You know, get out, you know, imagine, have sort of a, you know, they don't imagine things that haven't been done before, so you can go out and actually test real time. Right. And, and it's, it's creating candidates or imagining scenarios with a rationale that can be articulated rather than at the moment, selecting from a library and seeing which work.

Yeah, that's right.

Dr Philip Ball: How, how much in system like this, how important do you think it is that we be able to articulate the current, the modes

David Cox: of reasoning that are being used? Do you think that's, for us, it's very important. So, um, you know, often say, you know, in our. Our customers all see all our customers are basically the largest companies in the world.

So all the fortune 500, you know, that that's who we serve. Um, they all, I've talked to a lot of their CIOs and CTO. It's not that they know that AI is coming to transform their business. But the interesting thing is, and then how it's, it's really, it's not just a technology transformation. It's also a people transformation because if I was going to have an impact on them in the near and long-term, they were going to have to change how they work, right.

People are going to have to incorporate AI into the workflows of what they do. And for that to happen, the technology has to meet them halfway. So it has to be able to understand why the system is telling them to do things that, that falls into the sort of broad category of explainability, you know, like it has to be able to explain to you, like, yes, I, you know, I recommend you do this treatment, you know, the doctor's going to say, okay, fine.

You going to do that treatment? I don't believe that. Hey, I recommend you doing this treatment because you know, 43% of patients that were like this, you know, I have this property and this property, you know, had this, you know, this treatment was effective. This other treatment was not effective. They were actually able to articulate or say like, you know, we, you know, here's some examples of patients where we did and didn't do that.

And we're going to show you this, um, that's much more compelling, uh, for the system. You know, a doctor can then use that for decision support rather than just treating it as this like mysterious black box. And that's true, not just in healthcare, but in every application where. You know, if they ask, system's telling you to do something, you want to be able to understand why it's telling you to do that.

And I think that's very important, but th there's an interesting, uh, divergence is happening in the field right now around explainability, where some groups are, um, building systems that kind of come up with. Hawk explanations of the reasons why the system made the decision. And that's kind of interesting because it's sort of like, you know what, my, I, you know, I'm at home here at crash in the kitchen.

I come out of the kitchen, uh, the Cree jars broken on the floors, cookies everywhere. And my daughter's there. Now. She will give me an explanation about what happened and. You know, human understandable explanation, but it's not necessarily true. So there's this idea of building systems like we're in this interesting stage now where we need to build AI systems that can not only give us explanations for why they made the decisions, but they have to actually be true.

They have to be the real explanations. And I think one thing we mentioned earlier, causal inference is actually really important. And also symbolic reasoning are both actually really interesting paths to real explainability because. You know, if you want to understand why something happened or why decisions made, if you've been explained the causes that's, that's the language we want to understand.

At in and these mock systems the same, same way. It's like, you've got, I can show you a series of symbolic operations that said like, okay, the system did this, this did this. And it's operating on symbols, which is what we fundamentally operate on at our high level of consciousness. That's going to drive us a lot closer.

So in many ways I think solving the underlying problems that we need to solve are going to naturally give rise to interpretable explainable sort of AI systems. And I think. It's imperative that we have that explainability for us to be able to use and get the fruits of these AI systems

Dr Philip Ball: finally, whereas sort of agreed that probably the robot apocalypse isn't going to happen in the near future.

What, what are the real things we should

David Cox: be concerned about in developing this? Yeah, I think we're going to have a near term sort of awkward adolescent period in some industries of applying AI. Like we have to be careful about that AI. Uh, so, so AI that's deployed in an irresponsible or a clumsy way and IBM is that.

Deeply concerned about this. Like, oh our whole, we have E you're hugely invested in the trust of our customers. So trust is a really important pillar of what we do. So we want to make sure that the systems are robust. We want to make sure the systems are secure. I want to make sure the systems are explainable, cause that's going to engender trust.

And I think we're going, unfortunately, we're going to see lots of like, kind of wild west applications of AI that are just sloppy in that. So for instance, there was a court. Yeah. A system judicial system in Florida was using machine learning to make decisions about parole. And they claimed it was going to be on bias because they blinded the system to race variables.

But of course the system was just learning from data that, uh, you know, like the biases were in society. So if you baked decisions based on past data and in some ways this big data era it's, it's sort of guaranteed to give you back what the past was. It can't be forward looking, right. It takes the past and it gives it back to you.

So I think. That's something to watch. Like that's actually, you know, people talk about Skynet and Terminator and conscious machines. We're nowhere near any of that, but we are very close to verus and very awkward and clumsy misapplications of today's AI technology. People do also worry about job displacement and, and, and, and, um, you know, people being replaced by AI.

I don't think that I wouldn't say that's a misplaced concern, but. History has told us, and we actually have a project in the lab with Eric , he's written a couple of books on sort of AI and labor markets. We actually have ingested a huge number of job ads, uh, even going back to the 19th century and use it, use AI methods to ingest the ads and that looked at jobs over time and see how they've changed.

And the bottom line is that very few jobs go away, jobs change and the skills. So the distribution of skills that we use to do that. Change over time. So, you know, like, uh, uh, you know, an administrative assistance, you know, at one point typing was a very important skill. It's not like we don't have administrative assistants anymore.

They just have like the actual typing and dictation. Isn't something that you're do anymore has transformed. Now they do email, they do other things. Um, you look at at banking, you know, you might say like, oh my gosh, automated teller machines, surely that wiped out the population of bank tellers. Like surely there are no more bank tellers.

Right. We automated. Nothing's further than the truth. There's actually more bank tellers now than there were before. ATM's because banks transformed into, from a mechanical job of dealing out money to a customer service role. So I think we, we it's, it's not something we should fear. AI is not going to replace us, is going to augment us and then augmenting us.

It's going to change the distribution of skill. So it's going to be something that's going to evolve over time. Jobs, jobs will change, but they will not like we're going to suddenly. You know, as surplus to requirement where we're going to, we're going to see transformations. Those are the things we have to watch out for because you know, different jobs will be impacted at different extents.

We need to think seriously about how we rescale peoples that they, that nobody gets hurt. Um, but I think it fits into the same theme of, uh, the concerns are less about emergent apocalypse. The concerns are much more about how do we responsibly proceed in a deliberate way that helps them. And I think that's kind of, that's our position.

That's, that's what I believe is going to happen. Like the dangerous will be irresponsibility, um, more than, you know, something getting away from us and becoming dangerous. Well, that's fantastic. Thank you so much

Dr Philip Ball: is not going to replace us, David. Well, that's a relief. I was being flippant about the robot apocalypse that we're always being warned about by futurologists, you know, where Skynet and its terminators eradicate the humans, but it's node fear. Once we start thinking about creating artificial beings like Frankenstein's creature, Victor Frankenstein remember refused his creatures demand to make a female partner because he worried that the creatures might meet.

And then. Uh,

David Cox: race of devils will be propagated upon the earth who might make the very existence of the species of man, a condition precarious and full of terror.

Dr Philip Ball: Yep. As David said, the real fear is that we ourselves will use AI unwisely, trying to get it to do tasks. It isn't really capable of. This vision of a homegrown Frankenstein creature with an AI brain is in all honesty, fanciful we're no where near being able to put together something like that. But it seems certain that AI is going to carry on getting smarter, perhaps developing more human light powers of reasoning.

Almost certainly falling us sometimes if only for a short time that we're interacting with another sentience being the real question is how we're going to live alongside. What will be the social relations between human and machine. That's what I will explore in the final episode,

Continuity: How To Grow A Human: my Frankenstein Summer is written and presented by Dr. Philip Ball and directed and edited by Keith English. This show is brought to you by Aurra Studios. Listen to the full series on Apple podcasts or wherever you get your podcasts.